

1 IN THE UNITED STATES DISTRICT COURT
2 FOR THE NORTHERN DISTRICT OF CALIFORNIA

3
4 Apple Inc.,

5 Plaintiff,

6
7 v.
8

9 Telefonaktiebolaget LM Ericsson,
10 Inc. et al.,

11 Defendants.
12

Civil Action No.: 3:15-cv-00154

**DECLARATION OF THE INTERNET
ENGINEERING TASK FORCE, AN
ORGANIZED ACTIVITY OF THE
INTERNET SOCIETY**

13
14 I, Alexa Morris, based on my personal knowledge and information, hereby declare as follows:

- 15 1. I am Executive Director of the Internet Engineering Task Force ("IETF") and have held
16 this position since January 1, 2008.
- 17 2. Among my responsibilities as Executive Director, I act as the custodian of Internet-
18 Drafts for the IETF. The IETF is an organized activity of the Internet Society. The
19 Internet Society is a professional membership organization of Internet experts that
20 comments on policies and practices and oversees a number of other boards and task
21 forces dealing with network policy issues. Through my position at IETF, I have
22 personal knowledge of the facts stated herein.
- 23 3. I make this declaration based on my personal knowledge and information contained in
24 the business records of the IETF, or confirmation with other responsible IETF personnel
25 with such knowledge.
26
27
28

- 1 4. It is the regular practice of the IETF to publish Internet-Drafts and make them available
2 to the public on its website at www.ietf.org. The IETF maintains copies of Internet-
3 Drafts in the ordinary course of its regularly conducted activities. An Internet-Draft is a
4 working document of the IETF, its areas, and its working groups. Other groups may
5 also distribute working documents as Internet-Drafts. During the development of a
6 specification, draft versions of the document are made available for informal review and
7 comment by placing them in the IETF's Internet-Drafts directory. This makes an
8 evolving working document readily available to a wide audience, facilitating the process
9 of review and revision. See <http://www.ietf.org/id-info/>.
10
- 11 5. Attachment A hereto lists two Internet-Drafts, true and correct copies of which are
12 included as Exhibits 1-2 on the accompanying CD.
- 13 6. I personally reviewed the documents included as Exhibits 1-2 on the accompanying CD.
- 14 7. I hereby certify, in accordance with the requirements of Federal Rule of Evidence 902,
15 that Exhibit 1-2 on the accompanying CD constitute a record of regularly conducted
16 business activity which was (A) made at or near the time of the occurrence of the matters
17 set forth by, or from information transmitted by, a person with knowledge of those
18 matters; (B) kept in the course of the regularly conducted activity; and (C) made by the
19 regularly conducted activity as a regular practice.
20
- 21 8. Based on a search of IETF records, I have determined that the IETF maintained a true
22 and correct copy of Exhibit 1 bearing the Bates number IETF_000081-IETF_000101 on
23 the accompanying CD, titled "SSH Transport Layer Protocol," submitted by T. Ylonen
24 et al. to the IETF, and dated September 2002, in the ordinary course of its regularly
25 conducted activities.
26
27
28

1 9. Based on a search of IETF records and the IETF's course of conduct in publishing
2 Internet-Drafts, I have also determined that Exhibit 1 on the accompanying CD was
3 published on the IETF website (www.ietf.org) at least as of September 2002 and was
4 reasonably accessible to the public, and was disseminated or otherwise available to the
5 extent that persons interested and ordinarily skilled in the subject matter or art exercising
6 reasonable diligence could have located it.
7

8 10. Based on a search of IETF records, I have determined that the IETF maintained a true
9 and correct copy of Exhibit 2 bearing the Bates numbers IETF_000102-IETF_000109 on
10 the accompanying CD, titled "Key Derivation for Authentication, Integrity, and
11 Privacy," submitted by Marc Horowitz to the IETF, and dated August 1998, in the
12 ordinary course of its regularly conducted activities.
13

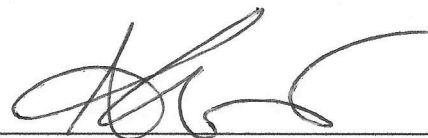
14 11. Based on a search of IETF records and the IETF's course of conduct in publishing
15 Internet-Drafts, I have also determined that Exhibit 2 on the accompanying CD was
16 published on the IETF website (www.ietf.org) at least as of August 1998 and was
17 reasonably accessible to the public, and was disseminated or otherwise available to the
18 extent that persons interested and ordinarily skilled in the subject matter or art exercising
19 reasonable diligence could have located it.
20

21 Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of
22 perjury under the laws of the United States of America that the foregoing is true and correct and
23 that the foregoing is based upon personal knowledge and information and is believed to be true.
24

25
26 Date:

October 18, 2015

By:



Alexa Morris

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

ATTACHMENT A

Authenticated Documents

Exhibit No.	Title	File Name	Publication Date
1	"SSH Transport Layer Protocol" by T. Ylonen et. al.,	draft-ietf-secsh-userauth-18.	September 2002
2	"Key Derivation for Authentication, Integrity, and Privacy" by Marc Horowitz	draft-horowitz-key-derivation-02	August 1998

- 1 4. It is the regular practice of the IETF to publish Internet-Drafts and make them available
2 to the public on its website at www.ietf.org. The IETF maintains copies of Internet-
3 Drafts in the ordinary course of its regularly conducted activities. An Internet-Draft is a
4 working document of the IETF, its areas, and its working groups. Other groups may
5 also distribute working documents as Internet-Drafts. During the development of a
6 specification, draft versions of the document are made available for informal review and
7 comment by placing them in the IETF's Internet-Drafts directory. This makes an
8 evolving working document readily available to a wide audience, facilitating the process
9 of review and revision. See <http://www.ietf.org/id-info/>.
10
11 5. Attachment A hereto lists two Internet-Drafts, true and correct copies of which are
12 included as Exhibits 1-2 on the accompanying CD.
13
14 6. I personally reviewed the documents included as Exhibits 1-2 on the accompanying CD.
15
16 7. I hereby certify, in accordance with the requirements of Federal Rule of Evidence 902,
17 that Exhibit 1-2 on the accompanying CD constitute a record of regularly conducted
18 business activity which was (A) made at or near the time of the occurrence of the matters
set forth by, or from information transmitted by, a person with knowledge of those

Network Working Group
Ylonen
Internet-Draft
Corp
Expires: March 2, 2003
Ed.
Inc
2002

T.
SSH Communications Security
D. Moffat,
Sun Microsystems,
September

SSH Authentication Protocol
draft-ietf-secsh-userauth-18.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 2, 2003.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

SSH is a protocol for secure remote login and other secure network services over an insecure network. This document describes the SSH authentication protocol framework and public key, password, and host-based client authentication methods. Additional authentication

methods are described in separate documents. The SSH authentication protocol runs on top of the SSH transport layer protocol and provides a single authenticated tunnel for the SSH connection protocol.

Ylonen & Moffat
1]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

Table of Contents

1.	Contributors	3
2.	Introduction	3
3.	Conventions Used in This Document	3
3.1	The Authentication Protocol Framework	3
3.1.1	Authentication Requests	4
3.1.2	Responses to Authentication Requests	5
3.1.3	The "none" Authentication Request	6
3.1.4	Completion of User Authentication	6
3.1.5	Banner Message	7
3.2	Authentication Protocol Message Numbers	7
3.3	Public Key Authentication Method: publickey	8
3.4	Password Authentication Method: password	10
3.5	Host-Based Authentication: hostbased	11
4.	Security Considerations	12
	Normative	13

13	Informative
14	Authors' Addresses
15	Intellectual Property and Copyright Statements

Ylonen & Moffat
2]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

1. Contributors

The major original contributors of this document were: Tatu Ylonen, Tero Kivinen, Timo J. Rinne, Sami Lehtinen (all of SSH Communications Security Corp), and Markku-Juhani O. Saarinen (University of

Jyvaskyla)

The document editor is: Darren.Moffat@Sun.COM. Comments on this internet draft should be sent to the IETF SECSH working group, details at: <http://ietf.org/html.charters/secsh-charter.html>

2. Introduction

The SSH authentication protocol is a general-purpose user authentication protocol. It is intended to be run over the SSH transport layer protocol [SSH-TRANS]. This protocol assumes that the underlying protocols provide integrity and confidentiality protection.

This document should be read only after reading the SSH architecture document [SSH-ARCH]. This document freely uses terminology and notation from the architecture document without reference or further explanation.

The service name for this protocol is "ssh-userauth".

When this protocol starts, it receives the session identifier from the lower-level protocol (this is the exchange hash H from the first key exchange). The session identifier uniquely identifies this session and is suitable for signing in order to prove ownership of a private key. This protocol also needs to know whether the lower-level protocol provides confidentiality protection.

3. Conventions Used in This Document

The keywords "MUST", "MUST NOT", "REQUIRED", "SHOULD", "SHOULD NOT", and "MAY" that appear in this document are to be interpreted as described in [RFC2119]

The used data types and terminology are specified in the architecture document [SSH-ARCH]

The architecture document also discusses the algorithm naming conventions that MUST be used with the SSH protocols.

3.1 The Authentication Protocol Framework

The server drives the authentication by telling the client which

Ylonen & Moffat
3]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

authentication methods can be used to continue the exchange at any given time. The client has the freedom to try the methods listed by the server in any order. This gives the server complete control over the authentication process if desired, but also gives enough flexibility for the client to use the methods it supports or that are most convenient for the user, when multiple methods are offered by the server.

Authentication methods are identified by their name, as defined in [SSH-ARCH]. The "none" method is reserved, and MUST NOT be listed as supported. However, it MAY be sent by the client. The server MUST always reject this request, unless the client is to be allowed in without any authentication, in which case the server MUST accept this request. The main purpose of sending this request is to get the list of supported methods from the server.

The server SHOULD have a timeout for authentication, and disconnect if the authentication has not been accepted within the timeout period. The RECOMMENDED timeout period is 10 minutes. Additionally, the implementation SHOULD limit the number of failed authentication attempts a client may perform in a single session (the RECOMMENDED limit is 20 attempts). If the threshold is exceeded, the server SHOULD disconnect.

3.1.1 Authentication Requests

All authentication requests MUST use the following message format. Only the first few fields are defined; the remaining fields depend on the authentication method.

byte	SSH_MSG_USERAUTH_REQUEST
string	user name (in ISO-10646 UTF-8 encoding [RFC2279])

string service name (in US-ASCII)
string method name (US-ASCII)
The rest of the packet is method-specific.

The user name and service are repeated in every new authentication attempt, and MAY change. The server implementation MUST carefully check them in every message, and MUST flush any accumulated authentication states if they change. If it is unable to flush some authentication state, it MUST disconnect if the user or service name changes.

The service name specifies the service to start after authentication.

There may be several different authenticated services provided. If the requested service is not available, the server MAY disconnect immediately or at any later time. Sending a proper disconnect message is RECOMMENDED. In any case, if the service does not exist,

Ylonen & Moffat
4]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

authentication MUST NOT be accepted.

If the requested user does not exist, the server MAY disconnect, or MAY send a bogus list of acceptable authentication methods, but never accept any. This makes it possible for the server to avoid disclosing information on which accounts exist. In any case, if the user does not exist, the authentication request MUST NOT be accepted.

While there is usually little point for clients to send requests that the server does not list as acceptable, sending such requests is not an error, and the server SHOULD simply reject requests that it does not recognize.

An authentication request MAY result in a further exchange of messages. All such messages depend on the authentication method

used, and the client MAY at any time continue with a new SSH_MSG_USERAUTH_REQUEST message, in which case the server MUST abandon the previous authentication attempt and continue with the new one.

3.1.2 Responses to Authentication Requests

If the server rejects the authentication request, it MUST respond with the following:

byte	SSH_MSG_USERAUTH_FAILURE
string	authentications that can continue
boolean	partial success

"Authentications that can continue" is a comma-separated list of authentication method names that may productively continue the authentication dialog.

It is RECOMMENDED that servers only include those methods in the list that are actually useful. However, it is not illegal to include methods that cannot be used to authenticate the user.

Already successfully completed authentications SHOULD NOT be included in the list, unless they really should be performed again for some reason.

"Partial success" MUST be TRUE if the authentication request to which this is a response was successful. It MUST be FALSE if the request was not successfully processed.

When the server accepts authentication, it MUST respond with the following:

Ylonen & Moffat
5]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

byte	SSH_MSG_USERAUTH_SUCCESS
------	--------------------------

Note that this is not sent after each step in a multi-method

authentication sequence, but only when the authentication is complete.

The client MAY send several authentication requests without waiting for responses from previous requests. The server MUST process each request completely and acknowledge any failed requests with a SSH_MSG_USERAUTH_FAILURE message before processing the next request.

A request that results in further exchange of messages will be aborted by a second request. It is not possible to send a second request without waiting for a response from the server, if the first request will result in further exchange of messages. No SSH_MSG_USERAUTH_FAILURE message will be sent for the aborted method.

SSH_MSG_USERAUTH_SUCCESS MUST be sent only once. When SSH_MSG_USERAUTH_SUCCESS has been sent, any further authentication requests received after that SHOULD be silently ignored.

Any non-authentication messages sent by the client after the request that resulted in SSH_MSG_USERAUTH_SUCCESS being sent MUST be passed to the service being run on top of this protocol. Such messages can be identified by their message numbers (see Section Message Numbers (Section 3.2)).

3.1.3 The "none" Authentication Request

A client may request a list of authentication methods that may continue by using the "none" authentication method.

If no authentication at all is needed for the user, the server MUST return SSH_MSG_USERAUTH_SUCCESS. Otherwise, the server MUST return SSH_MSG_USERAUTH_FAILURE and MAY return with it a list of authentication methods that can continue.

This method MUST NOT be listed as supported by the server.

3.1.4 Completion of User Authentication

Authentication is complete when the server has responded with SSH_MSG_USERAUTH_SUCCESS; all authentication related messages received after sending this message SHOULD be silently ignored.

After sending SSH_MSG_USERAUTH_SUCCESS, the server starts the requested service.

3.1.5 Banner Message

In some jurisdictions, sending a warning message before authentication may be relevant for getting legal protection. Many UNIX machines, for example, normally display text from `/etc/issue`, or use "tcp wrappers" or similar software to display a banner before issuing a login prompt.

The SSH server may send a `SSH_MSG_USERAUTH_BANNER` message at any time before authentication is successful. This message contains text to be displayed to the client user before authentication is attempted. The format is as follows:

byte	<code>SSH_MSG_USERAUTH_BANNER</code>
string	message (ISO-10646 UTF-8)
string	language tag (as defined in [RFC3066])

The client SHOULD by default display the message on the screen. However, since the message is likely to be sent for every login attempt, and since some client software will need to open a separate window for this warning, the client software may allow the user to explicitly disable the display of banners from the server. The message may consist of multiple lines.

If the message string is displayed, control character filtering discussed in [SSH-ARCH] SHOULD be used to avoid attacks by sending terminal control characters.

3.2 Authentication Protocol Message Numbers

All message numbers used by this authentication protocol are in the range from 50 to 79, which is part of the range reserved for protocols running on top of the SSH transport layer protocol.

Message numbers of 80 and higher are reserved for protocols running after this authentication protocol, so receiving one of them before

authentication is complete is an error, to which the server MUST respond by disconnecting (preferably with a proper disconnect message sent first to ease troubleshooting).

After successful authentication, such messages are passed to the higher-level service.

These are the general authentication message codes:

```
#define SSH_MSG_USERAUTH_REQUEST      50
#define SSH_MSG_USERAUTH_FAILURE      51
#define SSH_MSG_USERAUTH_SUCCESS      52
```

Ylonen & Moffat
7]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

```
#define SSH_MSG_USERAUTH_BANNER      53
```

In addition to the above, there is a range of message numbers (60..79) reserved for method-specific messages. These messages are only sent by the server (client sends only SSH_MSG_USERAUTH_REQUEST messages). Different authentication methods reuse the same message numbers.

3.3 Public Key Authentication Method: publickey

The only REQUIRED authentication method is public key authentication.

All implementations MUST support this method; however, not all users

need to have public keys, and most local policies are not likely to require public key authentication for all users in the near future.

With this method, the possession of a private key serves as authentication. This method works by sending a signature created with a private key of the user. The server MUST check that the key is a valid authenticator for the user, and MUST check that the signature is valid. If both hold, the authentication request MUST be accepted; otherwise it MUST be rejected. (Note that the server MAY require additional authentications after successful authentication.)

Private keys are often stored in an encrypted form at the client host, and the user must supply a passphrase before the signature can

be generated. Even if they are not, the signing operation involves some expensive computation. To avoid unnecessary processing and user

interaction, the following message is provided for querying whether authentication using the key would be acceptable.

```
byte      SSH_MSG_USERAUTH_REQUEST
string    user name
string    service
string    "publickey"
boolean   FALSE
string    public key algorithm name
string    public key blob
```

Public key algorithms are defined in the transport layer specification [SSH-TRANS]. The public key blob may contain certificates.

Any public key algorithm may be offered for use in authentication. In particular, the list is not constrained by what was negotiated during key exchange. If the server does not support some algorithm, it MUST simply reject the request.

The server MUST respond to this message with either

Ylonen & Moffat
8]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

SSH_MSG_USERAUTH_FAILURE or with the following:

```
byte      SSH_MSG_USERAUTH_PK_OK
string    public key algorithm name from the request
string    public key blob from the request
```

To perform actual authentication, the client MAY then send a signature generated using the private key. The client MAY send the signature directly without first verifying whether the key is acceptable. The signature is sent using the following packet:

```
byte      SSH_MSG_USERAUTH_REQUEST
```

```

string    user name
string    service
string    "publickey"
boolean   TRUE
string    public key algorithm name
string    public key to be used for authentication
string    signature

```

Signature is a signature by the corresponding private key over the following data, in the following order:

```

string    session identifier
byte      SSH_MSG_USERAUTH_REQUEST
string    user name
string    service
string    "publickey"
boolean   TRUE
string    public key algorithm name
string    public key to be used for authentication

```

When the server receives this message, it **MUST** check whether the supplied key is acceptable for authentication, and if so, it **MUST** check whether the signature is correct.

If both checks succeed, this method is successful. Note that the server may require additional authentications. The server **MUST** respond with `SSH_MSG_USERAUTH_SUCCESS` (if no more authentications are needed), or `SSH_MSG_USERAUTH_FAILURE` (if the request failed, or more authentications are needed).

The following method-specific message numbers are used by the publickey authentication method.

```

/* Key-based */
#define SSH_MSG_USERAUTH_PK_OK          60

```

3.4 Password Authentication Method: password

Password authentication uses the following packets. Note that a server MAY request the user to change the password. All implementations SHOULD support password authentication.

byte	SSH_MSG_USERAUTH_REQUEST
string	user name
string	service
string	"password"
boolean	FALSE
string	plaintext password (ISO-10646 UTF-8)

Note that the password is encoded in ISO-10646 UTF-8. It is up to the server how it interprets the password and validates it against the password database. However, if the client reads the password in some other encoding (e.g., ISO 8859-1 (ISO Latin1)), it MUST convert the password to ISO-10646 UTF-8 before transmitting, and the server MUST convert the password to the encoding used on that system for passwords.

Note that even though the cleartext password is transmitted in the packet, the entire packet is encrypted by the transport layer. Both

the server and the client should check whether the underlying transport layer provides confidentiality (i.e., if encryption is being used). If no confidentiality is provided (none cipher), password authentication SHOULD be disabled. If there is no confidentiality or no MAC, password change SHOULD be disabled.

Normally, the server responds to this message with success or failure. However, if the password has expired the server SHOULD indicate this by responding with SSH_MSG_USERAUTH_PASSWD_CHANGEREQ. In anycase the server MUST NOT allow an expired password to be used for authentication.

byte	SSH_MSG_USERAUTH_PASSWD_CHANGEREQ
string	prompt (ISO-10646 UTF-8)
string	language tag (as defined in [RFC3066])

In this case, the client MAY continue with a different authentication

method, or request a new password from the user and retry password authentication using the following message. The client MAY also send

this message instead of the normal password authentication request without the server asking for it.

byte	SSH_MSG_USERAUTH_REQUEST
string	user name

string service

Ylonen & Moffat
10]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

```
string      "password"
boolean     TRUE
string      plaintext old password (ISO-10646 UTF-8)
string      plaintext new password (ISO-10646 UTF-8)
```

The server must reply to request message with SSH_MSG_USERAUTH_SUCCESS, SSH_MSG_USERAUTH_FAILURE, or another SSH_MSG_USERAUTH_PASSWD_CHANGEREQ. The meaning of these is as follows:

SSH_MSG_USERAUTH_SUCCESS The password has been changed, and authentication has been successfully completed.

SSH_MSG_USERAUTH_FAILURE with partial success The password has been changed, but more authentications are needed.

SSH_MSG_USERAUTH_FAILURE without partial success The password
has not been changed. Either password changing was not supported,
or the old password was bad. Note that if the server has already sent SSH_MSG_USERAUTH_PASSWD_CHANGEREQ, we know that it supports changing the password.

SSH_MSG_USERAUTH_CHANGEREQ The password was not changed because the new password was not acceptable (e.g. too easy to guess).

The following method-specific message numbers are used by the password authentication method.

```
#define SSH_MSG_USERAUTH_PASSWD_CHANGEREQ    60
```

3.5 Host-Based Authentication: hostbased

Some sites wish to allow authentication based on the host where the user is coming from, and the user name on the remote host. While this form of authentication is not suitable for high-security sites,

it can be very convenient in many environments. This form of authentication is OPTIONAL. When used, special care SHOULD be taken to prevent a regular user from obtaining the private host key.

The client requests this form of authentication by sending the following message. It is similar to the UNIX "rhosts" and "hosts.equiv" styles of authentication, except that the identity of the client host is checked more rigorously.

This method works by having the client send a signature created with the private key of the client host, which the server checks with that host's public key. Once the client host's identity is established,

Ylonen & Moffat
11]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

authorization (but no further authentication) is performed based on the user names on the server and the client, and the client host name.

byte	SSH_MSG_USERAUTH_REQUEST
string	user name
string	service
string	"hostbased"
string	public key algorithm for host key
string	public host key and certificates for client host
string	client host name (FQDN; US-ASCII)
string	user name on the client host (ISO-10646 UTF-8)
string	signature

Public key algorithm names for use in "public key algorithm for host

key" are defined in the transport layer specification. The "public host key for client host" may include certificates.

Signature is a signature with the private host key of the following data, in this order:

string	session identifier
byte	SSH_MSG_USERAUTH_REQUEST
string	user name
string	service


```
string    "hostbased"
string    public key algorithm for host key
string    public host key and certificates for client host
string    client host name (FQDN; US-ASCII)
string    user name on the client host(ISO-10646 UTF-8)
```

The server MUST verify that the host key actually belongs to the client host named in the message, that the given user on that host is

allowed to log in, and that the signature is a valid signature on the appropriate value by the given host key. The server MAY ignore the client user name, if it wants to authenticate only the client host.

It is RECOMMENDED that whenever possible, the server perform additional checks to verify that the network address obtained from the (untrusted) network matches the given client host name. This makes exploiting compromised host keys more difficult. Note that this may require special handling for connections coming through a firewall.

4. Security Considerations

The purpose of this protocol is to perform client user authentication. It assumed that this runs over a secure transport

Ylonen & Moffat
12]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

layer protocol, which has already authenticated the server machine, established an encrypted communications channel, and computed a unique session identifier for this session. The transport layer provides forward secrecy for password authentication and other methods that rely on secret data.

Full security considerations for this protocol are provided in Section 8 of [SSH-ARCH]

Normative

[SSH-ARCH]

Ylonen, T., "SSH Protocol Architecture", I-D
draft-ietf-architecture-15.txt, Oct 2003.

[SSH-TRANS]

Ylonen, T., "SSH Transport Layer Protocol", I-D
draft-ietf-transport-17.txt, Oct 2003.

[SSH-USERAUTH]

Ylonen, T., "SSH Authentication Protocol", I-D
draft-ietf-userauth-18.txt, Oct 2003.

[SSH-CONNECT]

Ylonen, T., "SSH Connection Protocol", I-D
draft-ietf-connect-18.txt, Oct 2003.

[SSH-NUMBERS]

Lehtinen, S. and D. Moffat, "SSH Protocol Assigned
Numbers", I-D draft-ietf-secsh-assignednumbers-05.txt,
Oct 2003.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.

Informative

[RFC3066] Alvestrand, H., "Tags for the Identification of
Languages", BCP 47, RFC 3066, January 2001.

[RFC2279] Yergeau, F., "UTF-8, a transformation format of ISO
10646", RFC 2279, January 1998.

Ylonen & Moffat
13]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

Authors' Addresses

Tatu Ylonen
SSH Communications Security Corp
Fredrikinkatu 42
HELSINKI FIN-00100

Finland

EMail: ylo@ssh.com

Darren J. Moffat (editor)
Sun Microsystems, Inc
17 Network Circle
Menlo Park 95025
USA

EMail: Darren.Moffat@Sun.COM

Ylonen & Moffat
14]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be

followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

Ylonen & Moffat
15]

Expires March 2, 2003

[Page

Internet-Draft
2002

SSH Authentication Protocol

September

This document and the information contained herein is provided on an

"AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Ylonen & Moffat
16]

Expires March 2, 2003

[Page

Network Working Group
Horowitz
<draft-horowitz-key-derivation-02.txt>
Inc.
Internet-Draft
1998

M.
Stonecast,
August,

Key Derivation for Authentication, Integrity, and Privacy

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.ietf.org (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Distribution of this memo is unlimited. Please send comments to the author.

Abstract

Recent advances in cryptography have made it desirable to use longer cryptographic keys, and to make more careful use of these keys. In particular, it is considered unwise by some cryptographers to use the same key for multiple purposes. Since most cryptographic-based systems perform a range of functions, such as authentication, key exchange, integrity, and encryption, it is desirable to use different cryptographic keys for these purposes.

This RFC does not define a particular protocol, but defines a set

of
cryptographic transformations for use with arbitrary network
protocols and block cryptographic algorithm.

Deriving Keys

In order to use multiple keys for different functions, there are
two
possibilities:

- Each protocol ``key'' contains multiple cryptographic keys. The
implementation would know how to break up the protocol ``key''
for
use by the underlying cryptographic routines.
- The protocol ``key'' is used to derive the cryptographic keys.
The implementation would perform this derivation before calling

Horowitz
1]

[Page

Internet Draft
1997

Key Derivation

March,

the underlying cryptographic routines.

In the first solution, the system has the opportunity to provide
separate keys for different functions. This has the advantage that
if one of these keys is broken, the others remain secret. However,
this comes at the cost of larger ``keys'' at the protocol layer.

In
addition, since these ``keys'' may be encrypted, compromising the
cryptographic key which is used to encrypt them compromises all the
component keys. Also, the not all ``keys'' are used for all
possible
functions. Some ``keys'', especially those derived from passwords,
are generated from limited amounts of entropy. Wasting some of
this
entropy on cryptographic keys which are never used is unwise.

The second solution uses keys derived from a base key to perform
cryptographic operations. By carefully specifying how this key is
used, all of the advantages of the first solution can be kept,
while
eliminating some disadvantages. In particular, the base key must

be

used only for generating the derived keys, and this derivation must be non-invertible and entropy-preserving. Given these

restrictions,

compromise of one derived key does not compromise the other subkeys.

Attack of the base key is limited, since it is only used for derivation, and is not exposed to any user data.

Since the derived key has as much entropy as the base keys (if the cryptosystem is good), password-derived keys have the full benefit of

all the entropy in the password.

To generate a derived key from a base key:

$$\text{Derived Key} = \text{DK}(\text{Base Key}, \text{Well-Known Constant})$$

where

$$\text{DK}(\text{Key}, \text{Constant}) = \text{k-truncate}(\text{E}(\text{Key}, \text{Constant}))$$

In this construction, $\text{E}(\text{Key}, \text{Plaintext})$ is a block cipher, Constant is a well-known constant defined by the protocol, and k-truncate truncates its argument by taking the first k bits; here, k is the key size of E.

If the output of E is shorter than k bits, then some entropy in the key will be lost. If the Constant is smaller than the block size of E, then it must be padded so it may be encrypted. If the Constant is larger than the block size, then it must be folded down to the block size to avoid chaining, which affects the distribution of entropy.

In any of these situations, a variation of the above construction is used, where the folded Constant is encrypted, and the resulting output is fed back into the encryption as necessary (the | indicates concatenation):

$$\begin{aligned} K1 &= \text{E}(\text{Key}, \text{n-fold}(\text{Constant})) \\ K2 &= \text{E}(\text{Key}, K1) \end{aligned}$$

K3 = E(Key, K2)
K4 = ...

DK(Key, Constant) = k-truncate(K1 | K2 | K3 | K4 ...)

n-fold is an algorithm which takes m input bits and ``stretches'' them to form n output bits with no loss of entropy, as described in [Blumenthal96]. In this document, n-fold is always used to produce n bits of output, where n is the block size of E.

If the size of the Constant is not equal to the block size of E, then the Constant must be n-folded to the block size of E. This string is used as input to E. If the block size of E is less than the key size, then the output from E is taken as input to a second invocation of E. This process is repeated until the number of bits accumulated is greater than or equal to the key size of E. When enough bits have been computed, the first k are taken as the derived key.

Since the derived key is the result of one or more encryptions in the base key, deriving the base key from the derived key is equivalent to determining the key from a very small number of plaintext/ciphertext pairs. Thus, this construction is as strong as the cryptosystem itself.

Deriving Keys from Passwords

When protecting information with a password or other user data, it is necessary to convert an arbitrary bit string into an encryption key.

In addition, it is sometimes desirable that the transformation from password to key be difficult to reverse. A simple variation on the construction in the prior section can be used:

Key = DK(k-fold>Password), Well-Known Constant)

k-fold is same algorithm as n-fold, used to fold the Password into the same number of bits as the key of E.

The k-fold algorithm is reversible, so recovery of the k-fold output is equivalent to recovery of Password. However, recovering the k-fold output is difficult for the same reason recovering the base key from a derived key is difficult.

Traditionally, the transformation from plaintext to ciphertext, or vice versa, is determined by the cryptographic algorithm and the key.

A simple way to think of derived keys is that the transformation is determined by the cryptographic algorithm, the constant, and the key.

For interoperability, the constants used to derive keys for different purposes must be specified in the protocol specification. Also, the endian order of the keys must be specified.

Horowitz
3]

[Page

Internet Draft
1997

Key Derivation

March,

The constants must not be specified on the wire, or else an attacker who determined one derived key could provide the associated constant and spoof data using that derived key, rather than the one the protocol designer intended.

Determining which parts of a protocol require their own constants is an issue for the designer of protocol using derived keys.

Security Considerations

This entire document deals with security considerations relating to the use of cryptography in network protocols.

Appendix

This Appendix quotes the n-fold algorithm from [Blumenthal96]. It is provided here as a convenience to the implementor. Sample vectors are also included. It should be noted that the sample vector in Appendix B.2 of the original paper appears to be incorrect. Two independent implementations from the specification (one in C by the author, and another in Scheme by Bill Sommerfeld) agree on a value different from that in [Blumenthal96].

We first define a primitive called n-folding, which takes a variable-length input block and produces a fixed-length output sequence. The intent is to give each input bit approximately equal weight in determining the value of each output bit. Note that whenever we need to treat a string of bytes as a number, the assumed representation is Big-Endian -- Most Significant Byte first.

To n-fold a number X, replicate the input value to a length that is the least common multiple of n and the length of X. Before each repetition, the input is rotated to the right by 13 bit positions. The successive n-bit chunks are added together using 1's-complement addition (that is, with end-around carry) to yield a n-bit result....

The result is the n-fold of X. Here are some sample vectors, in hexadecimal. For convenience, the inputs are ASCII encodings of strings.

```
64-fold("012345") =  
64-fold(303132333435) = be072631276b1955
```

```
56-fold("password") =  
56-fold(70617373776f7264) = 78a07b6caf85fa
```

```
64-fold("Rough Consensus, and Running Code") =  
64-fold(526f75676820436f6e73656e7375732c20616e642052756e  
6e696e6720436f6465) = bb6ed30870b7f0e0
```

Horowitz
4]

[Page

Internet Draft
1997

Key Derivation

March,

```
168-fold("password") =  
168-fold(70617373776f7264) =  
59e4a8ca7c0385c3c37b3f6d2000247cb6e6bd5b3e
```

```
192-fold("MASSACHVSETTS INSTITVTE OF TECHNOLOGY"  
192-fold(4d41535341434856534554545320494e5354495456544520  
4f4620544543484e4f4c4f4759) =  
db3b0d8f0b061e603282b308a50841229ad798fab9540c1b
```

Acknowledgements

I would like to thank Uri Blumenthal, Hugo Krawczyk, and Bill Sommerfeld for their contributions to this document.

References

[Blumenthal96] Blumenthal, U., "A Better Key Schedule for DES-Like Ciphers", Proceedings of PRAGOCRYPT '96, 1996.

Author's Address

Marc Horowitz
Stonecast, Inc.
108 Stow Road
Harvard, MA 01451

Phone: +1 978 456 9103
Email: marc@stonecast.net

Horowitz
5]

[Page